The distribution of class groups of function fields

Jeffrey D. Achter
j.achter@colostate.edu

Abstract

For any sufficiently general family of curves over a finite field $\mathbb{F}_q$ and any elementary abelian ℓ-group H with ℓ relatively prime to q, we give an explicit formula for the proportion of curves C for which $\mathrm{Jac}(C)[\ell](\mathbb{F}_q) \cong H$. In doing so, we prove a conjecture of Friedman and Washington.

In 1983, Cohen and Lenstra introduced heuristics [5] to explain statistical observations about class groups of imaginary quadratic fields. Their principle, although still unproven, remains an important source of guidance in number theory. A concrete application of their heuristics predicts that an abelian group occurs as a class group of an imaginary quadratic field with frequency inversely proportional to the size of its automorphism group.

Six years later Friedman and Washington [9] addressed the function field case. Fix a finite field $\mathbb{F}_q$ and an abelian ℓ-group H, where ℓ is an odd prime relatively prime to q. Friedman and Washington conjecture that H occurs as the ℓ-Sylow part of the divisor class group of function fields over $\mathbb{F}_q$ with frequency inversely proportional to $|\mathrm{Aut}(H)|$. As evidence for this, they prove that the uniform distribution of Frobenius automorphisms of curves of genus g in $\mathrm{GL}_{2g}(\mathbb{Z}_\ell)$ would imply their conjecture. (Of course, autoduality of the Jacobian means that these Frobenius elements are actually in $\mathrm{GSp}_{2g}(\mathbb{Z}_\ell)$; still, [9] entertains the hope that this distinction is immaterial to the problem.) Friedman and Washington observe that, since geometric equidistribution results seem within reach, their conjecture may well be tractable.

Taking advantage of recent progress in equidistribution, we prove a statement in the spirit of [9]. A special, yet typical, case of our main result says the following.

Let $\mathcal{C} \to \mathcal{M}/\mathbb{F}_q$ be a relative smooth, proper curve of genus g over a smooth, irreducible variety, and suppose that $\mathcal{C} \to \mathcal{M}$ has full ℓ monodromy. Let $\{\mathbb{F}_{q^{en}}\}$ be a tower of finite extensions of $\mathbb{F}_q$ which is cofinal in the collection of all finite extensions of $\mathbb{F}_q$. For a curve C, let $\mathrm{Jac}(C)[\ell](k)$ denote the k-rational ℓ-torsion subgroup of its Jacobian. We give an explicit formula for a number $\alpha(g,r)$ so that:

$$\lim_{n\to\infty} \frac{\left|\{x \in \mathcal{M}(\mathbb{F}_{q^{en}}) : \mathrm{Jac}(\mathcal{C}_x)[\ell](\mathbb{F}_{q^{en}}) \cong (\mathbb{Z}/\ell)^r\}\right|}{\left|\mathcal{M}(\mathbb{F}_{q^{en}})\right|} = \alpha(g,r).$$

Moreover, $\lim_{g\to\infty} \alpha(g,r)$ exists. In this way, we can formulate a version of our result which allows the genus of the curves in question to change, too. The term $\alpha(g,r)$ should be thought of as a sort of symplectic analogue of $|\mathrm{Aut}((\mathbb{Z}/\ell)^r)|^{-1}$.

This result lets us essentially prove the original conjecture of Friedman and Washington. (We will see in Section 4.5 that the heuristic they introduce, that statistics of GL_{2g} should track those of Sp_{2g}, is a reasonable approximation but not literally true.) Moreover, we significantly strengthen (Section 4.3) results of Cardon and Murty [3] on divisibility of class groups of quadratic function fields.

The family of curves $\mathcal{C} \to \mathcal{M}$ is a concrete device for enumerating function fields. The "full ℓ-monodromy" constraint ensures that, as far as ℓ-Sylow subgroups of class groups are concerned, the family behaves like a general one.

On one hand, the familiar moduli spaces ${}_N\mathcal{M}_g$ [7] of proper smooth curves of genus g equipped with principal Jacobi level N structure have this property. In fact, so do most versal families of curves [8]. In this sense, our main theorem describes a typical collection of function fields of genus g.

On the other hand, it's not hard to write down families of curves which *do not* have this property. Generally speaking, extra algebraic cycles on a family of curves force the ℓ-adic monodromy to lie in a proper subgroup of Sp_{2g}. As a specific caution, we mention that if $d > 2$ then the family of curves $C_{d,f} : y^d = f(x)$ does not have full monodromy. (In fact, for such curves, $\mathbb{Z}[\zeta_d] \hookrightarrow \mathrm{End}_{\overline{\mathbb{F}}_q} \mathrm{Jac}(C_{d,f})$. One suspects [21] that the ℓ-adic monodromy is a unitary group associated to $\mathbb{Q}(\zeta_d)$; at the very least, the monodromy group is contained in such a group.)

The first section of this paper collects results of Katz about equidistribution of Frobenius elements on ℓ-adic sheaves. The second section investigates the combinatorics of $\mathrm{Sp}_{2g}(\mathbb{F}_\ell)$ and related groups. The next section combines these results to make precise statements about the distribution of class groups of function fields. The paper concludes with a series of applications of these results, culminating in a proof of a modified Friedman-Washington conjecture.

I thank R. Pries for helpful comments on this paper.

1 ℓ-adic monodromy

As noted above, Friedman and Washington foresaw that good equidistribution theorems would allow one to prove Cohen-Lenstra type results for function fields. Here, we recall the precise statements we need. Our discussion follows section one of [1], which itself is a recapitulation of parts of chapter nine of [15].

Fix an odd prime ℓ. Let $\mathcal{O}_\lambda$ be the ring of integers in some finite extension of $\mathbb{Z}_\ell$, and let $\Lambda = \mathcal{O}_\lambda/\lambda^n$ for some n. Let $V = V_\Lambda$ be a free, rank $2g$ Λ-module equipped with a symplectic form $\langle\cdot,\cdot\rangle$. The group of symplectic similitudes of $(V, \langle\cdot,\cdot\rangle)$ is

$$\mathrm{GSp}(V,\langle\cdot,\cdot\rangle) = \{A \in \mathrm{GL}(V) | \exists\, \mathrm{mult}(A) \in \Lambda^\times : \forall v,w \in V, \langle Av, Aw\rangle = \mathrm{mult}(A)\langle v,w\rangle\} \cong \mathrm{GSp}_{2g}(\Lambda).$$

The "multiplicator" mult is a character of $\mathrm{GSp}(V, \langle\cdot,\cdot\rangle)$, and its kernel is the usual symplectic group $\mathrm{Sp}_{2g}(\Lambda)$. For $\xi \in \Lambda^\times$, let $\mathrm{GSp}_{2g}^{\xi}(\Lambda) = \mathrm{mult}^{-1}(\xi)$ be the set of symplectic similitudes with multiplier ξ; each GSp_{2g}^{ξ} is a torsor over Sp_{2g}. For $W \subset \mathrm{GSp}_{2g}(\Lambda)$, let $W^{\xi} = W \cap \mathrm{GSp}_{2g}^{\xi}(\Lambda)$.

Let $T \to \operatorname{Spec}\mathbb{Z}[1/\ell]$ be a connected normal scheme of finite type, often $\operatorname{Spec}\mathbb{Z}[1/\ell]$ itself. Let $\mathcal{M} \to T$ be a scheme with smooth geometrically irreducible fibers; let $\eta_{\mathcal{M}}$ be a generic point of $\mathcal{M}$. A local system $\mathcal{F}$ of symplectic Λ-modules of rank $2g$ is equivalent to a continuous representation $\rho_{\mathcal{F}} : \pi_1(\mathcal{M}, \overline{\eta}_{\mathcal{M}}) \to \mathrm{Aut}(\mathcal{F}_{\overline{\eta}_{\mathcal{M}}}) \cong \mathrm{GSp}_{2g}(\Lambda)$. We call the image of this representation the (arithmetic) monodromy group of $\mathcal{F}$.

Let k be a finite field and $t \in T(k)$. Then $\mathcal{M}_t$ is a k-scheme, and we distinguish the geometric fundamental group $\pi_1^{\mathrm{geom}}(\mathcal{M}_t) = \pi_1(\mathcal{M}_t \times \overline{k}) \subseteq \pi_1(\mathcal{M}_t)$. The Galois group of k is (canonically isomorphic to) the quotient $\pi_1(\mathcal{M}_t)/\pi_1^{\mathrm{geom}}(\mathcal{M}_t)$.

We will require our sheaves to have uniform geometric monodromy group $G_{\mathrm{geom}} \subseteq \mathrm{Sp}_{2g}(\Lambda)$, in the sense that for every finite field k and $t \in T(k)$, the image of $\rho_{\mathcal{F}}(\pi_1^{\mathrm{geom}}(\mathcal{M}_t))$ is G_{geom}. With this assumption the monodromy group G, the full image of $\rho_{\mathcal{F}}$, is contained in $\Lambda \cdot G_{\mathrm{geom}} \subseteq \mathrm{GSp}_{2g}(\Lambda)$. We let $\xi(k)$ denote the image of Fr_k, the canonical generator of $\pi_1(\operatorname{Spec} k)$, in $G/G_{\mathrm{geom}} \subseteq \Lambda^\times$.

If k is a finite field, then to a k-point $x \in \mathcal{M}(k)$ one may associate its (conjugacy class of) Frobenius $\mathrm{Fr}_{x/k}$ in $\pi_1(\mathcal{M})$. Via $\rho_{\mathcal{F}}$, the point x acts on $\mathcal{F}_{\overline{\eta}_{\mathcal{M}}}$. Katz shows that these Frobenius elements are equidistributed in the monodromy group.

Theorem 1.1 (Katz). *Suppose $\mathcal{F}$ has uniform geometric monodromy group G_{geom} and arithmetic monodromy group G. Let $W \subset G$ be stable under G-conjugation. There are effective constants $\delta(\mathcal{M}, \mathcal{F})$ and $A(\mathcal{M}/T)$ so that, if k is a finite field with $|k| > A(\mathcal{M}/T)$ and $t : \operatorname{Spec} k \to T$ is an inclusion, then*

$$\left| \frac{|\{x \in \mathcal{M}_t(k) : \rho_{\mathcal{F}}(\mathrm{Fr}_{x,k}) \in W\}|}{|\mathcal{M}_t(k)|} - \frac{|W^{\xi(k)}|}{|G^{\xi(k)}|} \right| < \epsilon(\mathcal{M}, \mathcal{F}, k) := \frac{\delta(\mathcal{M}, \mathcal{F})}{\sqrt{|k|}}.$$

Proof This is simply [15, 9.7.13]; see also [4, 4.1]. While the result holds for any sheaf with finite monodromy group, we will (almost; see Section 4.2 below) always work with subgroups of $\mathrm{GSp}_{2g}(\Lambda)$. □

Already, this deep theorem yields a method for computing the proportion of curves in a family for which the ℓ-Sylow subgroup of the class group is isomorphic to a given group. Indeed, let H be any finite abelian group annihilated by ℓ^e, and let $\pi : \mathcal{C} \to \mathcal{M}/T$ be a smooth, irreducible proper relative curve of genus $g \geq 1$. For any finite field k and $t \in T(k)$ we define

$$\beta(\mathcal{C} \to \mathcal{M}, t, \ell^e, H) = \frac{|\{x \in \mathcal{M}_t(k) : \mathrm{Jac}(\mathcal{C}_x)[\ell^e](k) \cong H\}|}{|\mathcal{M}(k)|}. \tag{1.2}$$

We will often assume the k-point $t : \operatorname{Spec} k \to T$ is fixed, and simply write $\beta(\mathcal{C} \to \mathcal{M}, k, \ell^e, H)$.

There is a sheaf $\mathcal{F} = \mathcal{F}_{\mathcal{C},\ell^e}$ of abelian groups on $\mathcal{M}$ whose fiber at a geometric point $\overline{x} \in \mathcal{M}$ is the ℓ^e-torsion of the Jacobian $\mathrm{Jac}(\mathcal{C}_{\overline{x}})[\ell^e]$; it may be alternatively defined by $\mathcal{F} = R^1\pi_!(\mathbb{Z}/\ell^e)$. Suppose

that this family has uniform geometric monodromy group $G_{\text{geom}} \subseteq \mathrm{Sp}_{2g}(\mathbb{Z}/\ell^e)$ and arithmetic monodromy group $G \subseteq \mathrm{GSp}_{2g}(\mathbb{Z}/\ell^e)$. Then Theorem 1.1 implies that, for any sufficiently large finite field k and fixed, suppressed k-point of T, we have

$$\left|\beta(\mathcal{C} \to \mathcal{M}, k, \ell^e, H) - \frac{\left|\{x \in G^{\xi(k)} : \ker(x - \mathrm{id}) \cong H\}\right|}{\left|G^{\xi(k)}\right|}\right| < \epsilon(\mathcal{M}, \mathcal{F}, k). \tag{1.3}$$

We will denote the right-hand term of this inequality by $\epsilon_{\mathcal{C}\to\mathcal{M}}(\ell^e, k)$. Note that for a fixed family of curves $\mathcal{C} \to \mathcal{M}$, as $|k| \to \infty$ we have $\epsilon_{\mathcal{C}\to\mathcal{M}}(\ell^e, k) \to 0$.

In the special case where $\ell H = 0$, the mod-ℓ monodromy group is the full symplectic group, and $\{k_n\}$ is a collection of finite fields equipped with maps to T; $\lim_{n\to\infty} |k_n| = \infty$; and, for $n \gg 0$, $|k_n| \equiv 1 \bmod \ell$; we have

$$\lim_{n\to\infty} \beta(\mathcal{C} \to \mathcal{M}, k_n, \ell, H) = \frac{\left|\{x \in \mathrm{Sp}_{2g}(\mathbb{F}_\ell) : \ker(x - \mathrm{id}) \cong H\}\right|}{\left|\mathrm{Sp}_{2g}(\mathbb{F}_\ell)\right|}. \tag{1.4}$$

In the next section we explain how to compute the right-hand side of (1.4).

We collect the diverse notation and assumptions of this section in the following:

Situation 1.5. We suppose that ℓ is a fixed, odd prime; T is a connected $\mathbb{Z}[1/\ell]$-scheme of finite type; $\mathcal{M} \to \mathcal{T}$ is a smooth scheme with geometrically irreducible fibers; $\mathcal{C} \to \mathcal{M}$ is a proper, smooth relative curve of genus g; $\mathcal{F}_{\ell^e}$ is the sheaf of ℓ^e-torsion on the Jacobian of $\mathcal{C}$; and $\mathcal{F}$ has uniform geometric monodromy group G_{geom} and arithmetic monodromy group G. We will say that $\mathcal{C} \to \mathcal{M}$ has full ℓ^e-monodromy if $G_{\text{geom}} = \mathrm{Sp}_{2g}(\mathbb{Z}/\ell^e)$. Additionally, $\{k_n\}$ is a collection of finite fields, each equipped with an inclusion $t_n : k_n \to T$, such that $\lim_{n\to\infty} |k_n| = \infty$, and we will often write $\mathcal{M}(k)$ for $\mathcal{M}_{t_n}(k)$.

While any sequence of finite fields is allowed, psychologically it seems to be easiest to think of either $\{\mathbb{F}_{q^{e_n}}\}$, a tower of extensions of a fixed finite field $\mathbb{F}_q$, or $\{\mathbb{F}_{p_n}\}$, a collection of finite fields of ever-larger prime order.

2 Matrices with given fixed space

Katz's work on mod-ℓ monodromy reduces the calculation of β to a calculation in a symplectic group $\mathrm{Sp}_{2g}(\mathbb{F}_\ell)$. In this section we go to some length to calculate precisely the proportion of elements with a certain behavior. We note that [9] avoids these difficulties in two ways. First, the authors compute in GL_n, where the relevant Lie theory is more transparent, rather than in Sp_{2g}. Second, they compute with elements of Mat_n, with the hope that if x is equidistributed in GL_n, then $x - \mathrm{id}$ is equidistributed in Mat_n. Unfortunately, these choices mean that the conjectural description of class group frequencies in [9] differs slightly from the actual frequencies; we take up this point in more detail in Section 4.5.

2.1 Symplectic matrices over finite fields

Fix a finite field $\mathbb{F}$ with ℓ elements, where ℓ is a power of an odd prime. (In our applications ℓ will itself be prime, but this assumption is not necessary for the present computation.)

Our main goal is a formula for

$$\alpha(g,r) \quad := \quad \frac{\left|\{x \in \mathrm{Sp}_{2g}(\mathbb{F}) : \ker(x - \mathrm{id}) \cong \mathbb{F}^r\}\right|}{\left|\mathrm{Sp}_{2g}(\mathbb{F})\right|}. \tag{2.1.1}$$

The method presented here should work for any of the classical families of finite groups of Lie type. Still, since it is the symplectic group which arises most naturally in questions about the typical function field, we have chosen to focus our efforts on groups of type C. The reader will notice our heavy reliance on the paper [20] of Springer and Steinberg.

Consistent with the notation introduced in the previous section, we view $\mathrm{Sp}_{2g}(\mathbb{F})$ as the group of automorphisms of a $2g$-dimensional $\mathbb{F}$-vector space V_g equipped with a symplectic form $\langle \cdot, \cdot \rangle_g$. Unless otherwise noted, an r-subspace of V_g means any subspace $W \subset V_g$ for which $(W, \langle \cdot, \cdot \rangle_g|_W) \cong (V_r, \langle \cdot, \cdot \rangle_r)$. We will need to isolate the subspace of V_g on which a given element $x \in \mathrm{Sp}(V_g)$ acts unipotently.

Lemma 2.1.2. *Suppose $x \in \mathrm{Sp}(V) \cong \mathrm{Sp}_{2g}(\mathbb{F})$. Then there are subspaces $E_1(x)$ and $E_1(x)^\perp$ such that $V \cong E_1(X) \oplus E_1(x)^\perp$; $x|_{E_1(x)}$ is unipotent; and $x - \mathrm{id}$ is invertible on $E_1(x)^\perp$.*

Proof We assume that $x - \mathrm{id}$ is not invertible, as otherwise the statement is trivial. Write $V_{\overline{\mathbb{F}}}$ as the direct sum of generalized eigenspaces for x, $V_{\overline{\mathbb{F}}} = \oplus V_{\overline{\mathbb{F}}}(\lambda)$ where $V_{\overline{\mathbb{F}}}(\lambda)$ is the kernel of $(x - \lambda \,\mathrm{id})^{2g}$ on $V_{\overline{\mathbb{F}}}$.

Suppose that $\lambda\mu \neq 1$. We will prove, by induction on $m + n$, that

$$\langle \ker(x - \lambda \,\mathrm{id})^m|_{V_{\overline{\mathbb{F}}}}, \ker(x - \mu \,\mathrm{id})^n|_{V_{\overline{\mathbb{F}}}} \rangle = 0.$$

For the base case $m = n = 1$, suppose $xu = \lambda u$ and $xv = \mu v$. Because x preserves the symplectic form, we have $\langle u, v \rangle = \langle xu, xv \rangle = \langle \lambda u, \mu v \rangle = \lambda\mu \langle u, v \rangle$. Since $\lambda\mu \neq 1$, this forces $\langle u, v \rangle = 0$.

We now treat the inductive step. Suppose that $u \in \ker(x - \lambda \,\mathrm{id})^m|_{V_{\overline{\mathbb{F}}}}$ and $v \in \ker(x - \mu \,\mathrm{id})^n|_{V_{\overline{\mathbb{F}}}}$. Without loss of generality, assume that $m \geq n \geq 1$. Then $xu = u' + \lambda u$, where $u' \in \ker(x - \lambda \,\mathrm{id})^{m-1}|_{V_{\overline{\mathbb{F}}}}$, and $xv = v' + \mu v$ for some $v' \in \ker(x - \mu \,\mathrm{id})^{n-1}|_{V_{\overline{\mathbb{F}}}}$. (If $n = 1$, this simply means that $v' = 0$.) We then have $\langle u, v \rangle = \langle xu, xv \rangle = \langle u' + \lambda u, v' + \mu v \rangle = \langle u', v' \rangle + \lambda \langle u, v' \rangle + \mu \langle u', v \rangle + \lambda\mu \langle u, v \rangle$. By the inductive hypothesis, the first three terms in the last expression vanish. This leaves us with $\langle u, v \rangle = \lambda\mu \langle u, v \rangle$; again, $\langle u, v \rangle = 0$.

This shows that, if $\lambda\mu \neq 1$, then $\langle V_{\overline{\mathbb{F}}}(\lambda), V_{\overline{\mathbb{F}}}(\mu) \rangle = 0$. Since the pairing $\langle \cdot, \cdot \rangle$ is nondegenerate on V, we conclude that the pairing $\langle \cdot, \cdot \rangle : V_{\overline{\mathbb{F}}}(\lambda) \times V_{\overline{\mathbb{F}}}(\lambda^{-1}) \to \overline{\mathbb{F}}$ is nondegenerate. In particular, $V_{\overline{\mathbb{F}}}(1)$ is self-dual under $\langle \cdot, \cdot \rangle$.

Now, the generalized eigenspace associated to 1 is defined over $\mathbb{F}$; therefore, its orthogonal complement is, too. Returning to the $\mathbb{F}$-vector space V, we find that $E_1(x) := \ker(x - \mathrm{id})^{2g} \subset V$ is a symplectic subspace of V. Therefore, there exists a canonical decomposition $V = E_1(x) \oplus E_1(x)^\perp$, where 1 is not an eigenvalue of the action of x on $E_1(x)^\perp$.$\square$

We define the following quantities associated to $\mathrm{Sp}_{2g}(\mathbb{F})$. Let $\nu(g)$ be the number of elements in $\mathrm{Sp}_{2g}(\mathbb{F})$; let $U(g)$ be the number of unipotent elements in $\mathrm{Sp}_{2g}(\mathbb{F})$; and let $S(g,r)$ be the number of r-subspaces of V_g. Let $\Phi(g)$ be the number of elements $x \in \mathrm{Sp}_{2g}$ for which $x - \mathrm{id}$ is invertible, and let $\phi(g) = \Phi(g)/\nu(g)$ be the proportion of symplectic matrices with this property. For convenience, we define $\Phi(0) = 1$.

Recall that $\alpha(g,r)$ is the proportion of elements $x \in \mathrm{Sp}_{2g}(\mathbb{F})$ for which $\ker(x - \mathrm{id}) \cong \mathbb{F}^r$. Let $U(g,r)$ be the number of *unipotent* elements u of $\mathrm{Sp}_{2g}(\mathbb{F})$ for which $\ker(u - \mathrm{id}) \cong \mathbb{F}^r$. These quantities enjoy the following relations.

Lemma 2.1.3. *With all notation as above, let* $\lambda(j) = \ell^{2j-1}(\ell^{2j} - 1)$. *Then* $\nu(g) = \prod_{j=1}^{g} \lambda(j)$; $S(g,r) = \nu(g)/(\nu(r)\nu(g-r))$; $U(j) = \ell^{2j^2}$;

$$\alpha(g,r) = \frac{1}{\nu(g)} \sum_{j=1}^{g} S(g,j)U(j,r)\Phi(g-j); \tag{2.1.4}$$

and

$$\Phi(g) = \nu(g) - \sum_{j=1}^{g} S(g,j)U(j)\Phi(g-j). \tag{2.1.5}$$

Proof The calculation of ν and S is standard geometric algebra [2, III.6]. One proves that the symplectic group acts simply transitively on symplectic bases for V_g, and that $\lambda(g)$ counts the number of symplectic pairs in V_g. A theorem of Steinberg ([14, 8.14] or [20]) says that the number of unipotent elements in a finite group G of Lie type is $\ell^{\dim G - \mathrm{rank}\, G}$. Therefore $U(g)$, the number of unipotent elements in $\mathrm{Sp}_{2g}(\mathbb{F})$, is ℓ^{2g^2}.

By Lemma 2.1.2, any $x \in \mathrm{Sp}_{2g}(\mathbb{F})$ determines a decomposition $V = E_1(x) \oplus E_1(x)^\perp$, where x acts unipotently on $E_1(x)$ and $(x - \mathrm{id})$ is invertible on $E_1(x)^\perp$. Therefore, any element of the symplectic group determines, and is determined by, the data of a subspace $W \subset V$; a unipotent element $u \in \mathrm{Sp}(W)$; and an element $y \in \mathrm{Sp}(W^\perp)$ for which $(y - \mathrm{id})$ is invertible. If x corresponds in this way to the triple (W, u, y), then $\ker(x - \mathrm{id}) = \ker(u - \mathrm{id})|_W$.

Equations (2.1.4) and (2.1.5) follow swiftly. The right-hand side of (2.1.4) enumerates all choices of data (W, u, y) where W is a j-subspace of V, u is a unipotent element of $\mathrm{Sp}(W)$ with $\ker(u - \mathrm{id}) \cong \mathbb{F}^r$, and $y \in \mathrm{Sp}(W^\perp)$ with $y - \mathrm{id}$ invertible, all normalized by the size of the symplectic group.

To calculate $\Phi(g)$ and thus derive (2.1.5), we simply subtract from $\nu(g)$ the number of symplectic elements with nontrivial unipotent part. We enumerate triples (W, u, y) as before, where W is a

positive-dimensional subspace of V. If $W \cong V_j$, then $U(j)$ counts the number of choices for u, while $\Phi(g-j)$ is, by definition, the number of choices for y.$\square$

Equation (2.1.4), combined with Proposition 2.1.6 below, allows the explicit computation of $\alpha(g,r)$ in any particular case. The results of this calculation for $g \leq 3$ are shown in Table 4.1.

Proposition 2.1.6. *The number of unipotent elements u in $\mathrm{Sp}_{2g}(\mathbb{F})$ such that $\ker(u-\mathrm{id}) \cong \mathbb{F}^r$ is*

$$U(g,r) = \nu(g) \sum_{\mathbf{d}:0<d_1\leq d_2\leq\cdots\leq d_r} \left(\ell^{\frac{1}{2}\left(\sum_i s_i^2 - \sum_i r_i^2 + \sum_{i \text{ even}} r_i\right)} \cdot \prod_{i \text{ odd}} \nu(r_i/2) \cdot \prod_{i \text{ even}} \nu_{\mathrm{Orth}}(r_i) \right)^{-1} \tag{2.1.7}$$

where the sum is over all partitions $\mathbf{d}$ of $\dim V_g$ into r parts such that odd parts occur with even multiplicity; $r_i = \left|\{j : d_j = i\}\right|$; $s_i = \sum_{j\geq i} r_i$; and

$$\nu_{\mathrm{Orth}}(n) = \begin{cases} \ell^{m^2} \prod_{i=1}^{m} (\ell^{2i}-1) & n = 2m+1 \\ \ell^{m^2-2m} \prod_{i=1}^{m} (\ell^{2i}-1) & n = 2m \end{cases}.$$

Proof Let $G = \mathrm{Sp}_{2g}$. We start by identifying the relevant $G(\overline{\mathbb{F}})$-conjugacy classes of unipotent elements. As in, say, [14, 6.20], let $\mathcal{U} = \mathcal{U}(G)$ be the unipotent variety of G; it parametrizes all unipotent elements of G. Similarly, let $\mathcal{N}$ be the nilpotent variety of $\mathfrak{g}$, the Lie algebra of G. The Cayley transform is a G-equivariant isomorphism

$$\begin{array}{ccc} \mathcal{U} & \longrightarrow & \mathcal{N} \\ x & \longmapsto & (1-x)(1+x)^{-1} \end{array}$$

Thus, it suffices to count those $y \in \mathcal{N}(\mathbb{F})$ with nullspace of rank r.

Happily, enumeration of nilpotent elements is a classical result. Moreover, the description makes it easy to pick out those with the appropriate rank. To give a nilpotent orbit in $\mathfrak{sl}_n$ is to describe its Jordan normal form; a similar classification exists for arbitrary Lie groups. We have the classical bijection ([6, 5.1.1], [14, 7.11]) between nilpotent orbits of $\mathfrak{g}$ and the partitions of $2g$ for which odd parts occur with even multiplicity. The dimension of the nullspace of an element in a nilpotent orbit corresponding to a given partition is the number of elements in that partition. Therefore, the desired (geometric) nilpotent orbits are represented by suitable partitions with exactly r pieces. Each of these conjugacy classes has a representative in $G(\mathbb{F})$, and the summation in equation (2.1.7) thus ranges over all $G(\overline{\mathbb{F}})$-conjugacy classes of unipotent elements x in $G(\mathbb{F})$ for which $\ker(x-\mathrm{id}) \cong \mathbb{F}^r$.

We now explain how these $G(\overline{\mathbb{F}})$ conjugacy classes behave over $G(\mathbb{F})$, and compute the isomorphism class of the centralizer (still in $G(\mathbb{F})$) of an element of such a conjugacy class. By doing so, we are able to compute the size of the relevant conjugacy class.

We proceed as in [20, IV.2]. Fix a geometric conjugacy class corresponding to a partition $\mathbf{d}$ of g, and let $I = I(\mathbf{d}) = \{i : i \text{ even and } r_i > 0\}$. Jordan factors corresponding to even members d_i

split into two conjugacy classes over $\mathbb{F}$. Therefore, to give a $G(\mathbb{F})$-conjugacy class inside the $G(\overline{\mathbb{F}})$ conjugacy class $\mathbf{d}$ is to give a map of sets $c : I \to \{-1, +1\}$.

Let u be a representative for the $G(\mathbb{F})$-conjugacy class corresponding to $\mathbf{d}$ and a choice of assignments c. A theorem of Springer and Steinberg [20, IV.2.26-8] computes the isomorphism class of the centralizer $Z = Z_G(u)$ in G. It is the semidirect product of a unipotent radical, R, and the centralizer C of a certain torus associated to u. (Note that [20] computes the connected component of the centralizer, and then later accounts for multiple components.) The dimension of the Lie algebra of R is $\frac{1}{2}\left(\sum_i s_i^2 - \sum_i r_i^2 + \sum_{i \text{ even}} r_i\right)$. The reductive group C is isomorphic to

$$\prod_{i \text{ odd}} \mathrm{Sp}_{r_i}(\mathbb{F}) \cdot \prod_{i \text{ even}} \mathrm{O}_{r_i}^{c(i)}(\mathbb{F}).$$

Here, if r_i is even then $\mathrm{O}_{r_i}^{+1}(\mathbb{F})$ denotes the rank r_i orthogonal group of Witt defect 0 over $\mathbb{F}$, while $\mathrm{O}_{r_i}^{-1}(\mathbb{F})$ is the orthogonal group of Witt defect 1. For odd r_i, $\mathrm{O}_{r_i}^{\pm 1}(\mathbb{F})$ is the (unique) orthogonal group of rank r_i.

Therefore, the size of the set of elements in $G(\mathbb{F})$ which belong to the $G(\overline{\mathbb{F}})$-conjugacy class represented by $\mathbf{d}$ is

$$\begin{aligned}\sum_{c: I \to \{-1,+1\}} \frac{\left|\mathrm{Sp}_{2g}(\mathbb{F})\right|}{|R| \prod_{i \text{ odd}} \left|\mathrm{Sp}_{r_i}(\mathbb{F})\right|} \prod_{i \in I} \left|\mathrm{O}_{r_i}^{c(i)}(\mathbb{F})\right|^{-1} &= \frac{\nu(g)}{|R| \cdot \prod_{i \text{ odd}} \nu(r_i/2)} \cdot \prod_{i \in I} \left(\left|\mathrm{O}_{r_i}^{(-1)}(\mathbb{F})\right|^{-1} + \left|\mathrm{O}_{r_i}^{(+1)}(\mathbb{F})\right|^{-1} \right) \\ &= \frac{\nu(g)}{\ell^{\frac{1}{2}\left(\sum_i s_i^2 - \sum_i r_i^2 + \sum_{i \text{ even}} r_i\right)} \prod_{i \text{ odd}} \nu(r_i/2)} \cdot \prod_{i \in I} \nu_{\mathrm{Orth}}(r_i)\end{aligned}$$

where

$$\nu_{\mathrm{Orth}}(n) = \begin{cases} \ell^{m^2} \prod_{i=1}^{m} (\ell^{2i} - 1) & n = 2m+1 \\ \ell^{m^2 - 2m} \prod_{i=1}^{m} (\ell^{2i} - 1) & n = 2m \end{cases}.$$

Note that $\nu_{\mathrm{Orth}}(2m+1)$ is simply the number of elements in $\mathrm{SO}_{2m+1}(\mathbb{F})$, while $\nu_{\mathrm{Orth}}(2m+1)$ is the harmonic mean of $\left|\mathrm{SO}_{2m}^{(-1)}(\mathbb{F})\right|$ and $\left|\mathrm{SO}_{2m}^{(+1)}(\mathbb{F})\right|$. By summing over suitable geometric conjugacy classes $\mathbf{d}$ we obtain equation (2.1.7).□

Lemma 2.1.8. *The limits*

$$\phi(\infty) := \lim_{g \to \infty} \phi(g) \text{ and } \alpha(\infty, r) := \lim_{g \to \infty} \alpha(g, r)$$

exist.

Proof By Lemma 2.1.3,

$$\begin{aligned}\phi(g) &= \frac{1}{\nu(g)}(\nu(g) - \sum_{j=1}^{g} S(g,j) U(j) \Phi(g-j)) \\ &= 1 - \sum_{j=1}^{g} \frac{U(j)\Phi(g-j)}{\nu(j)\nu(g-j)}.\end{aligned}$$

Now, $\Phi(g-j)$ is necessarily less than $\nu(g-j)$, while $U(j)/\nu(j) < \ell^{-j}$. Therefore, $\lim_{g\to\infty}\phi(g)$ exists. Similarly, consider

$$\alpha(g,r) = \sum_{j=1}^{g} \frac{U(j,r)\Phi(g-j)}{\nu(j)\nu(g-j)}.$$

Again, $U(j,r)/\nu(j) \leq U(j)/\nu(j) < \ell^{-j}$, so that $\lim_{g\to\infty}\alpha(g,r)$ converges.□

2.2 Unitary groups

The methods of Section 2.1 work for any family of classical Lie groups. Since unitary groups also come up in certain natural applications (see Section 4.4), we briefly indicate how the argument works for $\mathbf{U}_n$. Because $\mathbf{U}_n$ is a twist of GL_n, the details are actually somewhat simpler. We preserve all notation from Section 2.1, using the subscript $\mathbf{U}$ to denote the appropriate group.

So, let $\mathbf{U}_n$ denote the unitary group in n variables over $\mathbb{F}$. Implicit in this definition is a nontrivial involution σ of $\mathbb{F}$; let m be $\sqrt{\ell}$, the size of the fixed field of σ. The number of elements in $\mathbf{U}_n$ is

$$\nu_{\mathbf{U}}(n) = m^{\frac{1}{2}(n^2-n)}\prod_{i=1}^{n}(m^i-(-1)^i);$$

the number of unipotent elements is $U_{\mathbf{U}}(n) = m^{n^2-n}$; and $S_{\mathbf{U}}(n,r) = \nu_{\mathbf{U}}(n)/(\nu_{\mathbf{U}}(r)\nu_{\mathbf{U}}(n-r))$. Moreover, the number of unitary matrices for which 1 is not an eigenvalue is

$$\Phi_{\mathbf{U}}(n) = \nu_{\mathbf{U}}(n) - \sum_{j=1}^{n} S_{\mathbf{U}}(n,j)U_{\mathbf{U}}(j)\Phi_{\mathbf{U}}(n-j).$$

Since the unitary group is a form of the general linear group, unipotent classes are parametrized by (unrestricted) partitions of n. Moreover, $\mathbf{U}_n(\overline{\mathbb{F}})$ and $\mathbf{U}_n(\mathbb{F})$ conjugacy coincide. The centralizer of a unipotent element corresponding to the partition $0 < d_1 \leq \cdots \leq d_r$ of n is connected, and has

$$m^{\sum s_i^2 - \sum r_i^2}\prod \nu_{\mathbf{U}}(d_i)$$

elements. All other results of Section 2.1, including the existence of $\alpha_{\mathbf{U}}(\infty,r)$, carry over.

3 Class groups of families of curves

Let H be a finite abelian group with $\ell^e H = 0$. We would like to compute the chance that H is the ℓ-part of the class group of a function field. As "sample space" of function fields we choose the fibers of any relative curve $\mathcal{C} \to \mathcal{M}/T$ as in Situation 1.5 with full ℓ^e-monodromy.

In practice, general families of curves tend to have full ℓ^e-monodromy; see, for instance, the introduction to [8]. As a concrete example, fix a natural number $N \geq 3$ relatively prime to p and

consider ${}_N\mathcal{C}_g \to {}_N\mathcal{M}_g$, the universal curve of genus g with principal Jacobi structure of level N. By [7, 5.15-5.16], this family of curves has full ℓ^e monodromy. Indeed, any versal family of curves has full monodromy at most primes [1, 2.2]. We also expect (see Section 4.3) that a general family of hyperelliptic curves has full ℓ^e-monodromy.

The equidistribution results in the first section let us detect the occurrence of H in class groups of function fields. Recall (Equation (1.2)) that this is measured by $\beta(\mathcal{C} \to \mathcal{M}, k, \ell^e, H)$.

Theorem 3.1. *Let H be a finite abelian ℓ-group. As in Situation 1.5, let $\mathcal{C} \to \mathcal{M}/T$ be a relative curve with full ℓ^e-monodromy and let $\{k_n\}$ be a collection of finite fields. Suppose that for $n \gg 0$, $\xi(k_n) \equiv \xi \bmod \ell^e$. There exists an effective constant $\delta(\mathcal{C} \to \mathcal{M})$ so that, for n sufficiently large,*

$$\left|\beta(\mathcal{C} \to \mathcal{M}, k_n, \ell^e, H) - \alpha^{\xi}(g, H, \ell^e)\right| < \epsilon_{\mathcal{C}\to\mathcal{M}}(\ell^e, k_n) := \frac{\delta(\mathcal{C} \to \mathcal{M})}{\sqrt{|k_n|}},$$

and thus

$$\lim_{n\to\infty} \beta(\mathcal{C} \to \mathcal{M}, k_n, \ell^e, H) = \alpha^{\xi}(g, H, \ell^e),$$

where

$$\alpha^{\xi}(g, H, \ell^e) = \frac{\left|\{x \in \mathrm{GSp}_{2g}^{\xi}(\mathbb{Z}/\ell^e) : \ker(x - \mathrm{id}) \cong H\}\right|}{\left|\mathrm{Sp}_{2g}(\mathbb{Z}/\ell^e)\right|}.$$

In the special case where $e = 1$ and $\xi = 1$, this term is computed by Lemma 2.1.3 and Proposition 2.1.6.

Proof As above consider the lisse sheaf $\mathcal{F} = \mathcal{F}_{\mathcal{C},\ell^e}$ on $\mathcal{M}$ which associates, to each geometric point $\overline{x}$, the ℓ^e-torsion of the Jacobian of $\mathcal{C}_{\overline{x}}$. Let $x \in \mathcal{M}(k_n)$ be any point. The divisor class group of $\mathcal{C}_x$ is $\mathrm{Jac}(\mathcal{C}_x)(k_n)$, the k_n-rational points of the Picard variety. By definition, the ℓ^e-torsion of this group is the subgroup of $\mathcal{F}_x$ fixed by Fr_{x/k_n}. Thus, in the notation of the first section, $\mathrm{Jac}(\mathcal{C}_x)[\ell^e](k_n) \cong \ker(\rho_{\mathcal{F}}(\mathrm{Fr}_{x/k_n}) - \mathrm{id})$, and

$$\beta(\mathcal{C} \to \mathcal{M}, k_n, \ell^e, H) = \frac{\left|\{x \in \mathcal{M}(k_n) : \ker(\rho_{\mathcal{F}}(\mathrm{Fr}_{x/k_n}) - \mathrm{id}) \cong H\}\right|}{|\mathcal{M}(k_n)|}.$$

In general, Theorem 1.1 finishes the proof. For the special case where H is an elementary abelian ℓ-group and $|k_n| \equiv 1 \bmod \ell$, Section 2 provides an algorithm for computing the appropriate quantity.□

In some applications, it is useful to be able to consider a family of curves with unbounded genus. To employ our methods, we need the size of the field of constants to grow more swiftly than the error terms ϵ of (1.3).

Theorem 3.2. *Let H be the elementary abelian ℓ-group $(\mathbb{Z}/\ell)^r$. As in Situation 1.5, let $\{\mathcal{C}_n \to \mathcal{M}_n/T_n\}_{n\in\mathbb{N}}$ be a collection of relative smooth proper curves of genus g_n with full ℓ-monodromy, and let $\{k_n\}$ be*

a collection of finite fields, each equipped with $t_n : \operatorname{Spec} k_n \to T_n$. *Suppose that* $\lim_{n\to\infty} g_n = \infty$; $\lim_{n\to\infty} \epsilon_{\mathcal{C}_n\to\mathcal{M}_n}(\ell, k_n) = 0$; *and for* $n \gg 0$, $\xi(k_n) = 1$. *Then*

$$\lim_{n\to\infty} \beta(\mathcal{C}_n \to \mathcal{M}_n, k_n, \ell, H) = \alpha(\infty, r),$$

where $\alpha(\infty, r)$ *is computed in Lemma 2.1.8.*

Proof The analysis is the same as that in Theorem 3.1. For n sufficiently large that $\xi(k_n) = 1$, we have

$$|\beta(\mathcal{C}_n \to \mathcal{M}_n, k_n, \ell, H) - \alpha(g_n, r)| < \epsilon_{\mathcal{C}_n\to\mathcal{M}_n}(\ell, k_n).$$

By Lemma 2.1.8, $\lim_{n\to\infty} \alpha(g_n, r)$ exists, with limit $\alpha(\infty, r)$. By hypothesis, $\lim_{n\to\infty} \epsilon_{\mathcal{C}_n\to\mathcal{M}_n}(\ell, k_n) = 0$; the theorem then follows.□

As predicted in [9], the divisor class groups of curves satisfy a Cohen-Lenstra type result. Recent research also addresses the distribution of other ideal class groups of function fields [3, 11, 18]. These studies work with an explicit affine model for a family of curves. To ease notation somewhat, we work with a relative curve $\mathcal{C} \to \mathcal{M}/k_0$ over a fixed finite field k_0, and specify an affine model by introducing a nonempty collection of sections $S = \{\sigma_1, \cdots, \sigma_n : \mathcal{M} \to \mathcal{C}\}$ with disjoint image. Since we need to pass to extension fields to apply our main result, we assume each σ_i is defined over the base field, k_0. For a curve C and a nonempty finite set of points S, let $\mathcal{O}_{C,S} = \cap_{P\notin S}\mathcal{O}_P$ be the ring of functions regular outside S. Let $\operatorname{cl}(\mathcal{O}_{C,S})$ be the ideal class group of this Dedekind domain, and let $\operatorname{cl}(\mathcal{O}_{C,S})_\ell$ be the ℓ-Sylow part of that group.

The techniques of this paper don't yield exact formulae for the frequency with which a given group H occurs as $\operatorname{cl}(\mathcal{O}_{C,S})_\ell$. Still, we can at least give bounds for the occurrence of ℓ-Sylow subgroups of given rank; these bounds are nontrivial if the genus of C is larger than $|S|$. Let $\operatorname{rank}_\ell H = \dim_{\mathbb{Z}/\ell} H/\ell H$.

Corollary 3.3. *Let* $\mathcal{C} \to \mathcal{M}/k_0$ *be a smooth proper relative curve of genus* g *with full* ℓ*-monodromy. Let* S *be a nonempty finite set of sections* $\sigma : \mathcal{M} \to \mathcal{C}$ *with disjoint image inside* $\mathcal{C}$, *and let* $S_x = \cup_{\sigma\in S}\sigma(x)$. *Let* k *be a sufficiently large finite extension of* k_0 *with* $|k| \equiv 1 \bmod \ell$. *For any nonnegative integer* r,

$$\frac{|\{x \in \mathcal{M}(k) : \operatorname{rank}_\ell \operatorname{cl}(\mathcal{O}_{\mathcal{C}_x,S_x}) \le r\}|}{|\mathcal{M}(k)|} \ge \sum_{j=0}^{r} \phi(g, j) - \epsilon_{\mathcal{C}\to\mathcal{M}}(k, \ell), \tag{3.4}$$

while

$$\frac{|\{x \in \mathcal{M}(k) : \operatorname{rank}_\ell \operatorname{cl}(\mathcal{O}_{\mathcal{C}_x,S_x}) \ge r\}|}{|\mathcal{M}(k)|} \ge \sum_{j=r+|S|}^{g} \phi(g, j) - \epsilon_{\mathcal{C}\to\mathcal{M}}(k, \ell). \tag{3.5}$$

Proof Given Theorem 3.1, all that's necessary is to relate the ideal class group to the divisor class group. By a theorem of F. K. Schmidt [19, proposition 1], there is an exact sequence of groups

$$0 \longrightarrow \frac{\mathcal{D}(\mathcal{C}_x, S_x)^0}{\mathcal{P}(\mathcal{C}_x, S_x)} \longrightarrow \mathrm{Jac}(\mathcal{C}_x)(k) \longrightarrow \mathrm{cl}(\mathcal{O}_{\mathcal{C}_x, S_x}) \longrightarrow 0, \tag{3.6}$$

where $\frac{\mathcal{D}(\mathcal{C}_x, S_x)^0}{\mathcal{P}(\mathcal{C}_x, S_x)}$ is the class group of divisors of degree zero represented by divisor classes supported at S. (The sequence (3.6) is exact on the right because the sections σ are defined over k.) On one hand, this shows that the ℓ-rank of the ideal class group is no bigger than that of the the divisor class group. On the other hand, the ℓ-rank of the kernel of the surjection $\mathrm{Jac}(\mathcal{C}_x)(k) \to \mathrm{cl}(\mathcal{O}_{\mathcal{C}_x, S_x})$ is at most $|S|$. These two observations yield equations (3.4) and (3.5), respectively.□

4 Examples

We conclude by working out some examples of these considerations. Specifically, we show how Theorem 3.1 and its variants, in conjunction with the calculations in Section 2, let us recover results of [17]; justify a heuristic used in [12]; improve the main results of [3], and a special case of [18]; and discuss the conjecture of [9].

For the most part, we work over a fixed finite field $k \cong \mathbb{F}_q$. We often phrase our results in terms of $\alpha(g, r)$, which is computed by Lemma 2.1.3 and Proposition 2.1.6. Values of $\alpha(g, r)$ for $g \leq 3$ are shown in Table 4.1.

4.1 Elliptic curves, $q \equiv 1 \bmod \ell$

The ℓ-torsion of a random elliptic curve E is

$$E[\ell](k) \cong \begin{cases} \{1\} & \text{with probability close to } \frac{\ell^2-\ell-1}{\ell^2-1} \\ \mathbb{Z}/\ell & \text{with probability close to } \frac{1}{\ell} \\ (\mathbb{Z}/\ell)^2 & \text{with probability close to } \frac{1}{\ell(\ell^2-1)} \end{cases} \tag{4.1.1}$$

in the following sense.

Let $\mathcal{E} \to \mathcal{M}$ be a non-isotrivial family of elliptic curves, such as the Legendre family (with affine model) $y^2 = x(x-1)(x-\lambda)$ over the λ-line. Such a family is versal, and therefore [1, 2.2] has full ℓ-monodromy for almost all ℓ; fix one such ℓ. With a slight simplification of the notation of Equation (1.2), let

$$\beta_{\mathcal{E}\to\mathcal{M}}(k, r) = \frac{|\{x \in \mathcal{M}(k) : \mathcal{E}_x[\ell](k) \cong (\mathbb{Z}/\ell)^r\}|}{|\mathcal{M}(k)|}$$

be the proportion of elliptic curves in our family, defined over k, for which the ℓ-torsion subgroup is isomorphic to $(\mathbb{Z}/\ell)^r$. Suppose that $|k| \equiv 1 \bmod \ell$ and $|k|$ is sufficiently large. Then Theorem 3.1 says that

$$|\beta_{\mathcal{E}\to\mathcal{M}}(k, r) - \alpha(1, r)| \leq \epsilon_{\mathcal{E}\to\mathcal{M}}(\ell, k),$$

g	r	$\alpha(g,r)$
1	0	$\frac{\ell^2-\ell-1}{\ell^2-1}$
1	1	$\frac{1}{\ell}$
1	2	$\frac{1}{\ell(\ell^2-1)}$
2	0	$\frac{\ell^6-\ell^5-\ell^4+\ell+1}{(\ell^2-1)\left(\ell^4-1\right)}$
2	1	$\frac{\ell^3-\ell-1}{\ell^2(\ell^2-1)}$
2	2	$\frac{\ell^3-\ell-1}{\ell^2(\ell^2-1)^2}$
2	3	$\frac{1}{(\ell^2-1)\ell^4}$
2	4	$\frac{1}{\ell^4(\ell^2-1)\left(\ell^4-1\right)}$
3	0	$\frac{\ell^{12}-\ell^{11}-\ell^{10}+\ell^7+\ell^5+\ell^4-\ell^3-\ell-1}{(\ell^2-1)\left(\ell^4-1\right)(\ell^6-1)}$
3	1	$\frac{\ell^8-\ell^6+\ell^2-\ell^5+\ell-\ell^4+1}{\ell^3(\ell^2-1)\left(\ell^4-1\right)}$
3	2	$\frac{\ell^8-\ell^6+\ell^2-\ell^5+\ell-\ell^4+1}{\ell^3(\ell^2-1)^2\left(\ell^4-1\right)}$
3	3	$\frac{\ell^5-\ell^3-1}{\ell^7(\ell^2-1)^2}$
3	4	$\frac{\ell^5-\ell^3-1}{\ell^7(\ell^2-1)^2\left(\ell^4-1\right)}$
3	5	$\frac{1}{(\ell^2-1)\left(\ell^4-1\right)\ell^9}$
3	6	$\frac{1}{\ell^9(\ell^2-1)\left(\ell^4-1\right)(\ell^6-1)}$

Table 4.1: The proportion of symplectic matrices of dimension $2g$ with fixed subspace of exact dimension r, as computed in Lemma 2.1.3 and Proposition 2.1.6.

where the error term decays as $1/\sqrt{|k|}$, and $\alpha(1,r)$, defined in Equation (2.1.1), may be read off from the first section of Table 4.1.

4.2 Elliptic curves, $q \not\equiv 1 \bmod \ell$

In the situation $\mathcal{E} \to \mathcal{M}$ considered above, suppose that k is a large finite field for which $|k| \equiv \xi \not\equiv 1 \bmod \ell$. Again, Theorem 3.1 says that

$$\left|\beta_{\mathcal{E}\to\mathcal{M}}(k,r) - \alpha^{\xi}(1,r)\right| \leq \epsilon_{\mathcal{E}\to\mathcal{M}}(\ell,k), \tag{4.2.1}$$

where

$$\alpha^{\xi}(1,r) := \frac{\left|\{x \in \mathrm{GSp}_2^{\xi}(\mathbb{Z}/\ell) : \ker(x-\mathrm{id}) \cong (\mathbb{Z}/\ell)^r\}\right|}{|\mathrm{SL}_2(\mathbb{Z}/\ell)|}.$$

We have not computed $\alpha^{\xi}(g,r)$ in general, but it is not hard to compute $\alpha^{\xi}(1,r)$ directly (see also [1, 3.3]): if $\xi \neq 1$, then

$$\alpha^{\xi}(1,r) = \begin{cases} \frac{\ell-2}{\ell-1} & r=0 \\ \frac{1}{\ell-1} & r=1 \\ 0 & r=2 \end{cases}. \tag{4.2.2}$$

Note that this is compatible with the familiar result (use the Weil pairing) that if k has no ℓ^{th} root of unity, then an elliptic curve over k cannot have all its ℓ-torsion defined over k.

Taken together, Equations (4.1.1), (4.2.1) and (4.2.2), in the special case where k is a prime field $\mathbb{F}_p$, fully recover Theorem 1.14 of [17].

In [12], Gekeler studies the distribution of Frobenius elements of elliptic curves over $\mathbb{F}_p$, taken as elements of $\mathrm{GL}_2(\mathbb{Z}_\ell)$. Among other results, he computes the proportion of elements in $\mathrm{GL}_2(\mathbb{Z}/\ell^e)$ with given trace and determinant. (This is easier than the analogous question in $\mathrm{Sp}_{2g}(\mathbb{Z}/\ell^e)$, first because conjugacy and stable conjugacy coincide in GL_2, and second because of the severe constraints on Jordan blocks of 2×2 matrices.) Combining [12, 4.4] and Theorem 3.1 allows one to compute the proportion of elliptic curves with ℓ^e-torsion isomorphic to a given abelian ℓ-group H.

Moreover, we can justify a heuristic used in section 3 of [12]. There, it is asserted that if m and n are relatively prime, then for a fixed elliptic curve E/k, the actions of Frobenius on $E[m](\overline{k})$ and $E[n](\overline{k})$ are independent, at least if $m \cdot n$ is small relative to $\sqrt{|k|}$. Indeed, let $\mathcal{C} \to \mathcal{M} \to k_0$ be any relative curve with full mn-monodromy; for simplicity, assume that $|k_0| \equiv 1 \bmod mn$. Then Frobenius elements of Jacobians of curves $\mathcal{C}_x$ are, by Theorem 1.1, equidistributed in $\mathrm{Sp}_{2g}(\mathbb{Z}/mn)$. Since this latter group is isomorphic to $\mathrm{Sp}_{2g}(\mathbb{Z}/m) \oplus \mathrm{Sp}_{2g}(\mathbb{Z}/n)$, Gekeler's claim follows.

4.3 Quadratic function fields

Attention has recently turned to the explicit construction of ideal classes of given order in the class groups of quadratic function fields $\mathbb{F}_q(x, \sqrt{f(x)})$. Friesen computes both empirical [10] and

analytic [11] bounds for the chance that ℓ divides the class number of $\mathbb{F}_q(x)[y]/(y^2 - f(x))$, where f is a quartic polynomial. Cardon and Murty [3] show that there are at least $q^{d(\frac{1}{2}+\frac{1}{\ell})}$ imaginary quadratic extensions $K = \mathbb{F}_q(x, \sqrt{f(x)})$ of $\mathbb{F}_q(x)$ where $\deg f \leq d$ and the ideal class group of K has an element of prime order $\ell \geq 3$. While as q gets large this produces arbitrarily large families of quadratic function fields with class number divisible by ℓ, it is a vanishingly small proportion of all quadratic function fields.

We can use Corollary 3.3 to compute the *proportion* of quadratic function fields with class number divisible by ℓ, and thereby strengthen these results.

Suppose q is a power of an odd prime and that $q \equiv 1 \bmod \ell$. We let $k = \mathbb{F}_q$, and let $\{k_n\}$ be any collection of finite extensions of k with $\lim_{n\to\infty} |k_n| = \infty$.

Let $\mathcal{H}_d$ be the space of separable monic polynomials $f(x)$ of degree d. Over it lies $\mathcal{C}_d$, the curve with affine model $y^2 = f(x)$; it is a hyperelliptic curve of genus $\lfloor \frac{d-1}{2} \rfloor$. We work under the hypothesis that $\mathcal{C}_d \to \mathcal{H}_d$ has full ℓ-monodromy. For odd d, this is implied by unpublished work of J.K. Yu [15, 10.5.10]; we will treat the general case in a future work.

The function field of the curve with affine model $y^2 = f(x)$ is called an imaginary quadratic function field if $d = \deg f$ is odd, and a real quadratic function field otherwise. We address these cases separately.

If d is odd, then there is a single point "at infinity" in this affine model; the left-hand term of 3.6 is trivial, and the ideal class group of this ring is isomorphic to the $\mathbb{F}_q$-rational points of the Jacobian of the associated proper curve. We see that, for instance,

$$\lim_{n\to\infty} \frac{\left|\{f(x) \in k_n[x] : \deg f = d, f \text{ monic}, \ell | \mathrm{cl}(k_n[x, \sqrt{f(x)}])\}\right|}{|\{f(x) \in k_n[x] : \deg f = d, f \text{ monic}\}|}$$

is equal to

$$\lim_{n\to\infty} \frac{\left|\{f(x) \in k_n[x] : \deg f = d, f \text{ monic, separable}, \ell | \mathrm{cl}(k_n[x, \sqrt{f(x)}])\}\right|}{|\{f(x) \in k_n[x] : \deg f = d, f \text{ monic, separable}\}|},$$

since most polynomials are separable, which is in turn equal to

$$\lim_{n\to\infty} \frac{|\{x \in \mathcal{H}_d(k_n) : \ell || \mathrm{Jac}(\mathcal{C}_x)[\ell](k_n)|\}|}{|\mathcal{H}_d(k_n)|},$$

or $1 - \alpha(g, 0)$.

If d is even, then there are two points at infinity, and the regulator term in 3.6 is an abelian group on a single generator. Therefore, for any curve C/k_n with affine model $C^{\mathrm{aff}} : y^2 = f(x)$, we have

$$\mathrm{rank}_\ell(\mathrm{Jac}(C)[\ell](k_n)) \geq \mathrm{rank}_\ell(\mathrm{cl}(\mathcal{O}_{C^{\mathrm{aff}}})) \geq \mathrm{rank}_\ell(\mathrm{Jac}(C)[\ell](k_n)) - 1,$$

and the chance that ℓ divides the class group of the affine coordinate ring is bounded from below by

$$\sum_{r=2}^{g} \alpha(g,r) - \epsilon_{\mathcal{C}_d \to \mathcal{H}_d}(k_n, \ell).$$

Thus, we can significantly strengthen the main conclusion of [3]; as $n \to \infty$, there is an element of order ℓ in the class group of $k_n[x][y]/(y^2 - f(x))$ for a positive proportion of monic degree d polynomials $f(x) \in k_n[x]$.

Note that Theorem 3.2 allows one to make uniform statements about curves of the form $y^2 = f(x)$ as $\deg f \to \infty$, provided that the size of the base field grows sufficiently quickly.

4.4 Cyclic cubic fields

Pacelli [18] looks at curves $y^d = f(x)$, and obtains results (for general d) similar to those of [3] for $d = 2$. As mentioned in the introduction, the general family of curves (with affine model) $y^d = f(x)$ cannot have full mod ℓ monodromy, because of the extra automorphisms this family possesses. Still, by computing in the appropriate monodromy group one can calculate divisibility of class numbers for these families. We expect [21] that the monodromy group is a unitary group associated to $\mathbb{Q}(\zeta_d)$.

We take up these considerations in the special case where $d = 3$, the degree of f is 4, and 3 is invertible in the base field. Let $\mathcal{C} \to \mathcal{P}$ be the family of curves with affine model $y^3 = f(x)$, where f ranges over all separable polynomials of degree 4. Each fiber $\mathcal{C}_x$ has genus 3. Moreover, since there is an obvious action of a cyclic group of order 3 on $\mathcal{C}$, the Jacobian $\mathrm{Jac}(\mathcal{C})$ admits an action by $\mathbb{Z}[\zeta_3]$. The action on the tangent space at the identity of any fiber has signature $(2,1)$, since actions of type $(3,0)$ are rigid. Therefore, under the Torelli map, $\mathcal{C} \to \mathcal{P}$ becomes identified with an open subset of the Picard modular variety associated to $\mathbb{Z}[\zeta_3]$. Using transcendental arguments [13] and the theory of compactification [16], one knows that for almost all ℓ, the full ℓ-adic monodromy group of this family is $G(\mathbb{Z}_\ell)$, where G is the unitary group in three variables associated to $\mathbb{Z}[\zeta_3]$.

Suppose, then, that $\mathcal{C} \to \mathcal{P}$ has ℓ-monodromy group $G(\mathbb{Z}/\ell)$. If $\mathbb{Z}[\zeta_3]$ is inert at ℓ, then $G(\mathbb{Z}/\ell)$ is an example of the unitary groups studied in Section 2.2. In particular, we see that:

$$\mathrm{Jac}(y^3 = f(x))[\ell](k) \cong \begin{cases} \{1\} & \text{with probability close to } \frac{\ell\left(\ell^5-\ell^3-1\right)}{(\ell+1)(\ell^2-1)(\ell^3+1)} \\ (\mathbb{Z}/\ell)^2 & \text{with probability close to } \frac{\ell^5-\ell^2+\ell^4-\ell-1}{\ell^2(\ell+1)^2(\ell^2-1)} \\ (\mathbb{Z}/\ell)^4 & \text{with probability close to } \frac{\ell^3+\ell^2-1}{(\ell+1)^2(\ell^2-1)\ell^3} \\ (\mathbb{Z}/\ell)^6 & \text{with probability close to } \frac{1}{\ell^3(\ell+1)(\ell^2-1)(\ell^3+1)} \end{cases}.$$

(If $\mathbb{Z}[\zeta_d]$ splits at ℓ, then $G(\mathbb{Z}_\ell)$ is isomorphic to a general linear group, and a similar, but easier, calculation applies.)

4.5 The Friedman-Washington conjecture

Let $\mathcal{C} \to \mathcal{M}$ be a family of curves of genus g with full ℓ-monodromy; this corresponds to any suitably general family of curves. Let k be a large finite field, say with $|k| \equiv 1 \bmod \ell$. Then Theorem 3.1 says that the proportion of fibers $\mathcal{C}_x$, for $x \in \mathcal{M}(k)$, with $\mathrm{Jac}(\mathcal{C}_x)[\ell](k) \cong (\mathbb{Z}/\ell)^r$ is $\alpha(g,r)$; see Table 4.1 for the first few values of $\alpha(g,r)$.

Friedman and Washington [9] give a conjectural description of the frequency with which a given abelian ℓ-group occurs as the ℓ-Sylow part of the divisor class group of a function field. While they formulate their conjecture in terms of hyperelliptic curves in order to preserve the analogy with the Cohen-Lenstra heuristics, all of their arguments depend on merely having a sufficiently general family of curves. Given our expectation (Section 4.3) that hyperelliptic curves behave, in terms of ℓ-monodromy, like general curves, we compare the predictions of [9] to the results of Theorem 3.1 for $\mathcal{C} \to \mathcal{M}$ with full monodromy.

To facilitate this comparison we estimate the chance that the ℓ-part of the class group of a curve is trivial. Let $\phi_{\mathrm{GL}}(n)$ denote the proportion of elements $x \in \mathrm{GL}_n$ for which $x - \mathrm{id}$ is invertible. It is shown that for large n $\phi_{\mathrm{GL}}(n)$ approaches

$$\widetilde{\phi}_{\mathrm{GL}}(n) := \prod_{j=1}^{n}(1 - \ell^{-j}).$$

In the special case of genus 2, Friedman and Washington predict that the proportion of curves with trivial ℓ-class group is (close to) $\widetilde{\phi}_{\mathrm{GL}}(4)$, while Theorem 3.1 says that this proportion is actually $\alpha(2,0) = \frac{\ell^6-\ell^5-\ell^4+\ell+1}{(\ell^2-1)(\ell^4-1)}$. The gap between $\widetilde{\phi}_{\mathrm{GL}}(4)$ and $\alpha(2,0)$ is of order $1/\ell^2$.

Now, [9, p.131] expresses the hope that this discrepancy disappears for large genus; unfortunately, this difference persists. The proportion of curves of (arbitrarily large) genus with trivial ℓ-group approaches (the well-defined limit; see Lemma 2.1.8) $\phi(\infty)$, which by [1, 3.3] is $1 - \ell/(\ell^2-1) + O(1/\ell^3)$. The difference between the conjectural estimate of [9] and the actual value remains of order $1/\ell^2$, even as the genus of the curves in question gets arbitrarily large.

References

[1] Jeffrey D. Achter and Joshua Holden. Notes on an analogue of the Fontaine-Mazur conjecture. *Journal de Théorie des Nombres de Bordeaux*, 15(3), 2003.

[2] E. Artin. *Geometric algebra*. John Wiley & Sons Inc., New York, 1988. Reprint of the 1957 original, A Wiley-Interscience Publication.

[3] David A. Cardon and M. Ram Murty. Exponents of class groups of quadratic function fields over finite fields. *Canad. Math. Bull.*, 44(4):398–407, 2001.

[4] Nick Chavdarov. The generic irreducibility of the numerator of the zeta function in a family of curves with large monodromy. *Duke Math. J.*, 87(1):151–180, 1997.

[5] H. Cohen and H. W. Lenstra, Jr. Heuristics on class groups of number fields. In *Number theory, Noordwijkerhout 1983 (Noordwijkerhout, 1983)*, volume 1068 of *Lecture Notes in Math.*, pages 33–62. Springer, Berlin, 1984.

[6] David H. Collingwood and William M. McGovern. *Nilpotent orbits in semisimple Lie algebras*. Van Nostrand Reinhold Mathematics Series. Van Nostrand Reinhold Co., New York, 1993.

[7] P. Deligne and D. Mumford. The irreducibility of the space of curves of given genus. *Inst. Hautes Études Sci. Publ. Math.*, (36):75–109, 1969.

[8] Torsten Ekedahl. The action of monodromy on torsion points of Jacobians. In *Arithmetic algebraic geometry (Texel, 1989)*, pages 41–49. Birkhäuser Boston, Boston, MA, 1991.

[9] Eduardo Friedman and Lawrence C. Washington. On the distribution of divisor class groups of curves over a finite field. In *Théorie des nombres (Quebec, PQ, 1987)*, pages 227–239. de Gruyter, Berlin, 1989.

[10] Christian Friesen. Class group frequencies of real quadratic function fields: the degree 4 case. *Math. Comp.*, 69(231):1213–1228, 2000.

[11] Christian Friesen. Bounds for frequencies of class groups of real quadratic genus 1 function fields. *Acta Arith.*, 96(4):313–331, 2001.

[12] Ernst-Ulrich Gekeler. Frobenius distributions of elliptic curves over finite prime fields. *Int. Math. Res. Not.*, (37):1999–2018, 2003.

[13] Rolf-Peter Holzapfel. *The ball and some Hilbert problems*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, 1995.

[14] James E. Humphreys. *Conjugacy classes in semisimple algebraic groups*, volume 43 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1995.

[15] Nicholas M. Katz and Peter Sarnak. *Random matrices, Frobenius eigenvalues, and monodromy*. American Mathematical Society, Providence, RI, 1999.

[16] Michael J. Larsen. Arithmetic compactification of some Shimura surfaces. In *The zeta functions of Picard modular surfaces*, pages 31–45. Univ. Montréal, Montreal, QC, 1992.

[17] H. W. Lenstra, Jr. Factoring integers with elliptic curves. *Ann. of Math. (2)*, 126(3):649–673, 1987.

[18] Allison M. Pacelli. Abelian subgroups of any order in class groups of global function fields. *J. Number Theory*, 106(1):26–49, 2004.

[19] Michael Rosen. S-units and S-class group in algebraic function fields. *J. Algebra*, 26:98–108, 1973.

[20] T. A. Springer and R. Steinberg. Conjugacy classes. In *Seminar on Algebraic Groups and Related Finite Groups (The Institute for Advanced Study, Princeton, N.J., 1968/69)*, Lecture Notes in Mathematics, Vol. 131, pages 167–266. Springer, Berlin, 1970.

[21] Yuri G. Zarhin. The endomorphism rings of Jacobians of cyclic covers of the projective line. *Math. Proc. Cambridge Philos. Soc.*, 136(2):257–267, 2004.